REMARKS/ARGUMENTS

Claims 1-14 are pending. Claims 1, 9, and 10 have been amended. New claims 15-20 have been added. No new matter has been introduced. Applicants believe the claims comply with 35 U.S.C. § 112.

Claims 1-14 stand rejected under 35 U.S.C. § 102(e) as being anticipated by Bridge (US 6,530,035).

The present application relates to storage systems, and in particular to storage system management in which failure boundaries are taken into consideration when assigning storage volumes. Specifically, failure boundaries are determined to ensure that replication storage volumes cross failure boundaries so that the impact of a failure event within a given failure boundary may be minimized.

Applicants respectfully submit that independent claim 1 as amended is novel and patentable over Bridge because, for instance, Bridge does not teach or suggest the additional limitation that the primary and secondary storage volumes utilize horizontal or vertical addressing. As described in paragraph [0031] and shown in Figs. 4 and 5, two addressing formats can be used to address the storage volumes and the addressing format may have an impact upon the failure boundaries utilized within the storage system. More specifically, paragraph [0031] recites:

> Next will be described two major addressing formats -- horizontal and vertical. The addressing format, as will be seen, impacts the manner in which failure boundaries are considered. Figure 4 illustrates horizontal addressing. As suggested by the name, addresses in horizontal addressing are assigned across the storage volumes.... As mentioned, preferably the groups will cross failure boundaries. The level of the failure boundary will be determined automatically using the system software and an appropriate policy, or the level may be determined by a system administrator. However determined, the level can consist of one error correction group 301, a controller pair 302, etc., as discussed above. (Emphasis added: ¶[0031])

As a result of the use of horizontal addressing, an impact is made upon how failure boundaries are selected. Alternatively, vertical addressing could also be used to generate a different physical arrangement of the storage volumes.

In contrast, *Bridge fails to teach or suggest the horizontal or vertical addressing of storage volumes.* In column 8, line 31 to column 9, line 4 and Figs. 2A-2B,

Bridge does describe the use of striping to divide extents into a variety of stripe units, allowing for the stripe units to be accessed in parallel to increase the I/O throughput of the storage system. He further describes the addressing of extents either sequentially (in reference to Fig. 2B) or in parallel (in reference to Fig. 2C) However, the extents only represent a "contiguous section of storage on a named drive," as described in column 6, lines 16-17 and do not represent an entire storage volume, as described in amended claim 1. No mention is made at all of the horizontal or vertical addressing of storage volumes.

Applicants respectfully submit that independent claim 10 as amended is novel and patentable over Bridge because, for instance, Bridge does not teach or suggest determining the failure boundary using error correction group and controller group information for the set of primary storage volumes and the set of replication storage volumes. The error correction groups include one or more of the logical volumes used within the storage system. As further stated in ¶[0029] of the present application, "[t]he failure boundary is used to designate that if one of the disk drives 121 in the ECC group 301 fails, all of the other logical volumes which belong to that group 301 will be impacted." Fig. 3 further illustrates this point in showing that the ECC groups are made up of volumes:

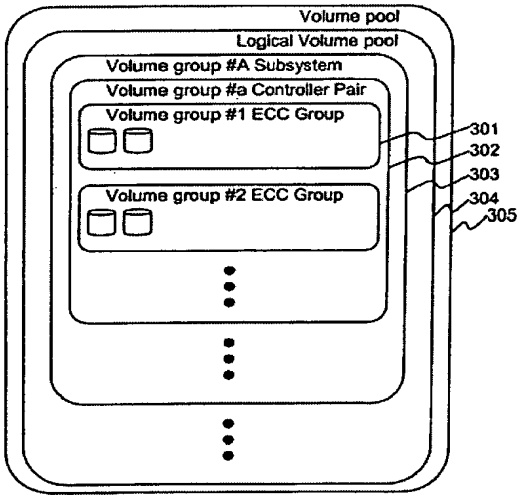

In contrast, the failure groups within Bridge are divided amongst disk drives, not logical volumes. As described in column 13, lines 32-56, each disk drive can be associated with a failure group where they share some common failure criteria. Bridge specifically states in regard to Fig. 6: "Disk drives 624, 626, 628, and 630 share a common

projected failure condition which results in these disk drives being assigned to the same failure group 622. Thus, disk drives are in the same failure group if there is a failure mode that could affect each of the disk drives and redundant data is maintained to protect against that failure." Furthermore, Fig. 6 illustrates three embodiments of the present invention in which disk drives are assigned as failure groups, instead of logical volumes:

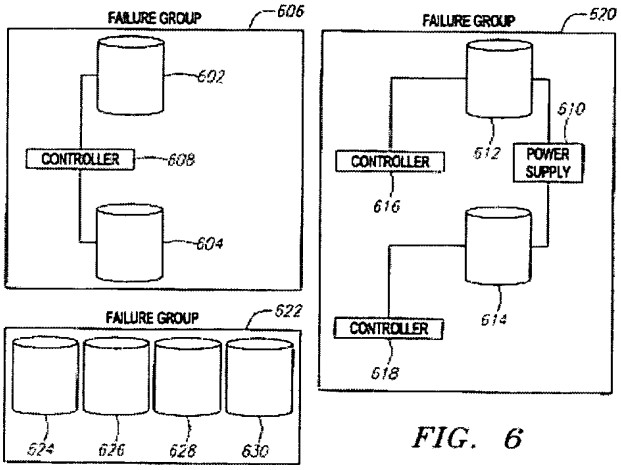

FIG. 6

Furthermore, the parity protection method utilized within Bridge utilizes the mirroring of data extents or contiguous storage sections on disks rather than separating failure boundaries based on error correction groups.

For at least the foregoing reasons, Applicants respectfully assert that independent claims 1 and 10, and dependent claims 2-9 and 11-14, are novel and patentable over Bridge.

New claims 15-20 depend from claims 1 and 10 and further recite determining a boundary of potential failure based on logical addresses and performing a replication process between the primary replication volumes and secondary storage volumes. These features are also absent from Bridge. Therefore, claims 15-20 are patentable over Bridge.

CONCLUSION

In view of the foregoing, Applicants believe all claims now pending in this Application are in condition for allowance. The issuance of a formal Notice of Allowance at an early date is respectfully requested.

If the Examiner believes a telephone conference would expedite prosecution of this application, please telephone the undersigned at 650-326-2400.

Respectfully submitted,

Chun-Pok Leung
Reg. No. 41,405

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, Eighth Floor
San Francisco, California 94111-3834
Tel: 650-326-2400
Fax: 415-576-0300
Attachments
RL:dlw
60746496 v1